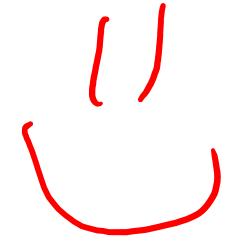


Announcements

Next week

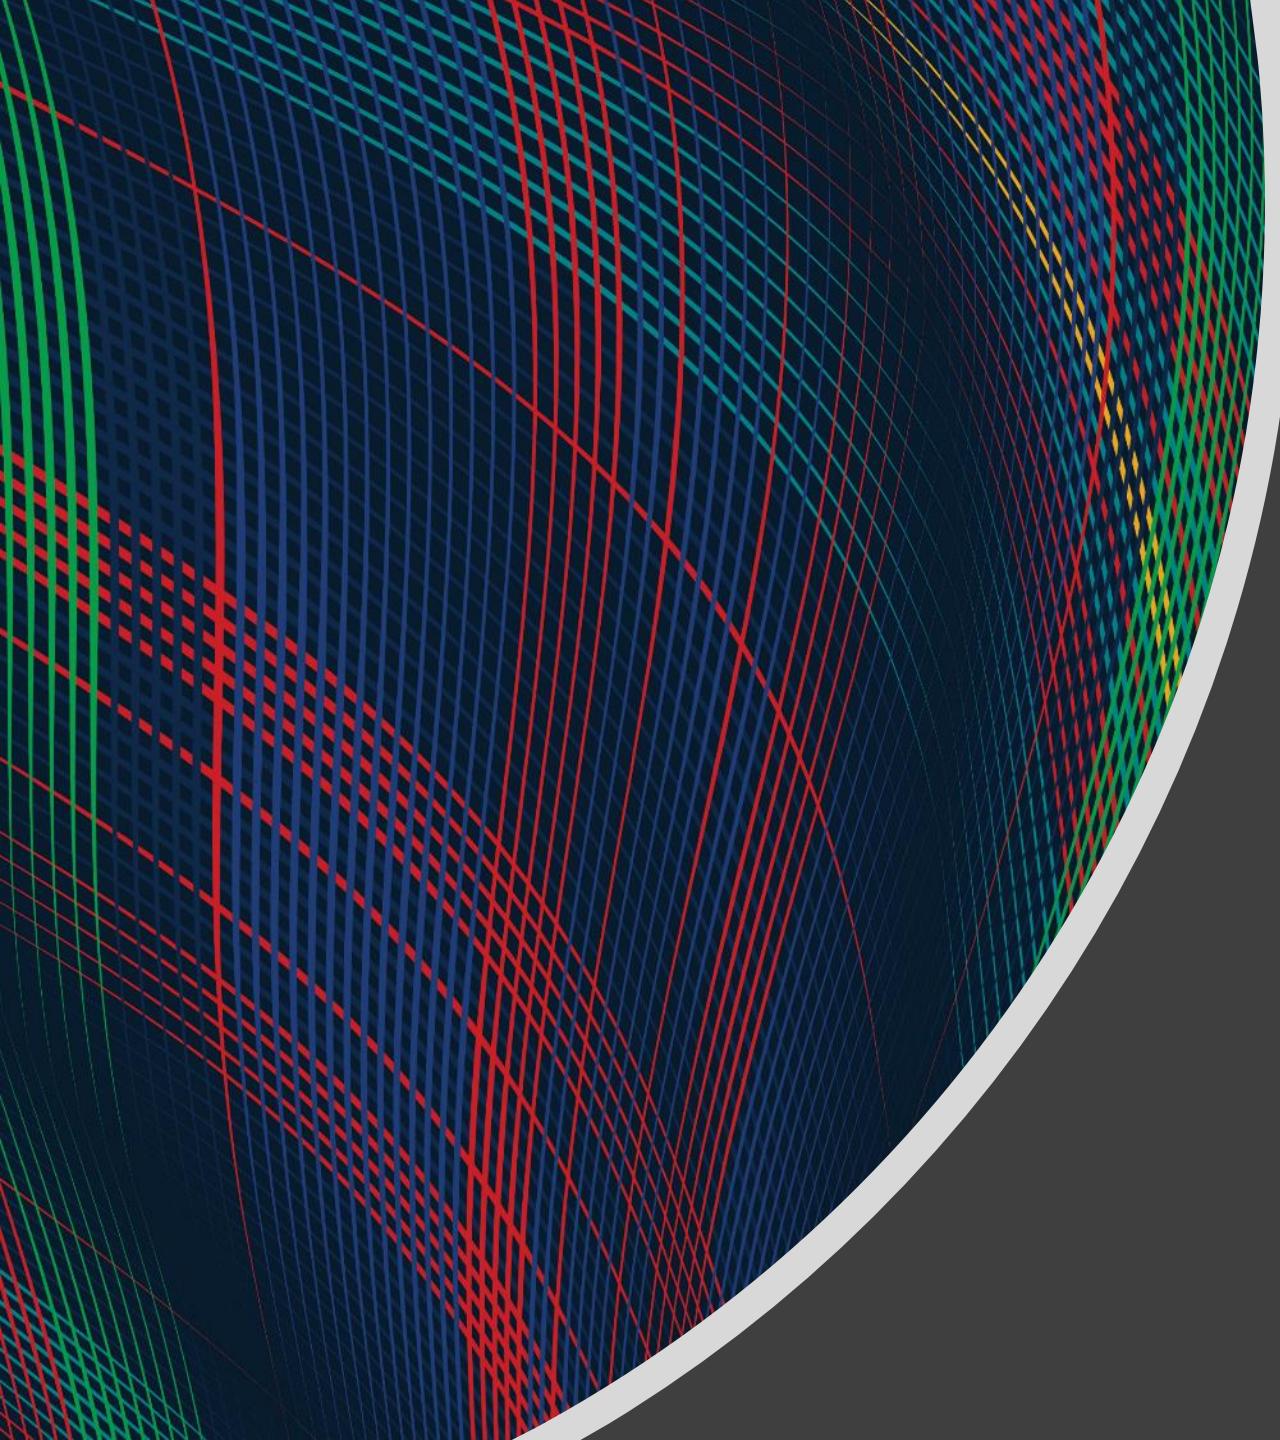


HW3

- Out late tonight
- Due Sat 10/9

Quizzes

- Mon 10/4, last 15 min. of class (calculus, optimization, Lagrange)
- Mon 10/11, last 15 min. of class (probability, statistics)



Mathematical Foundations for Machine Learning

Statistics

Instructor: Pat Virtue

Plan

Recitation

Trick coin

- Finding the best parameters
 - Among five different buckets
 - Plotting the log likelihood vs q

Today

Maximum likelihood estimation

Two applications of Bayes rule

Conditional likelihood

Expectation

Variance

Likelihood

Likelihood: The probability (or density) of random variable Y taking on value y given the distribution parameters, θ .

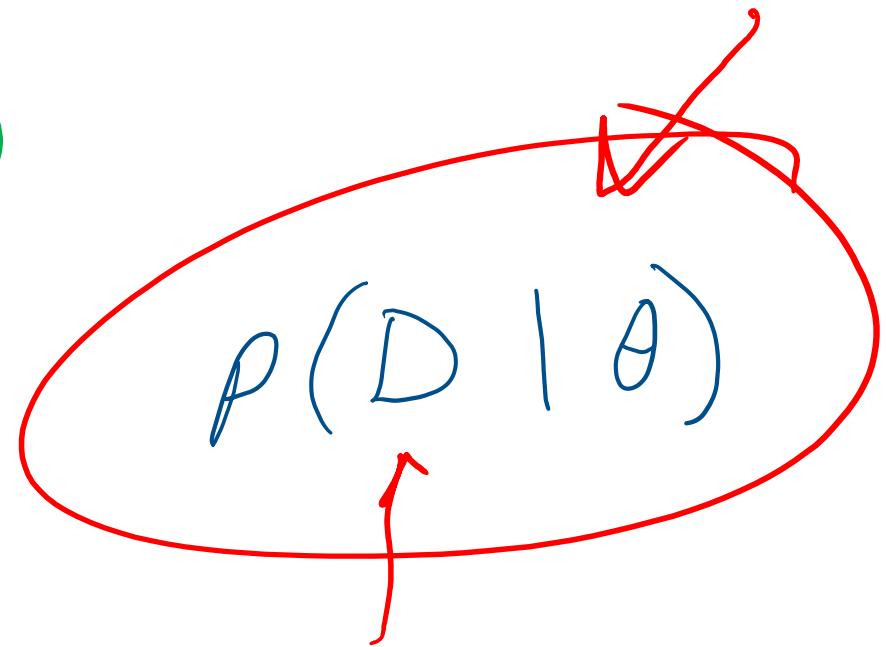
$$P(Y = y | \theta)$$

$$P(y | \theta)$$

$$P(y_j | \theta)$$

$$P_\theta(y)$$

Additional
notation



Likelihood and i.i.d

$$D = \{y^{(1)}, y^{(2)}, y^{(3)}\}$$

Likelihood: The probability (or density) of random variable Y taking on value y given the distribution parameters, θ .

i.i.d.: Independent and identically distributed

$$P(Y^{(1)} = y^{(1)}, Y^{(2)} = y^{(2)}, Y^{(3)} = y^{(3)} | \theta^{(1)}, \theta^{(2)}, \theta^{(3)})$$

identical



$$P(Y = y^{(1)}, Y = y^{(2)}, Y = y^{(3)} | \theta)$$

independent



$$= P(Y = y^{(1)} | \theta) P(Y = y^{(2)} | \theta) P(Y = y^{(3)} | \theta) = \prod_{i=1}^N P(Y^{(i)} | \theta)$$

$$P(D | \theta)$$

Bernoulli Likelihood

$$p(\mathcal{D} | \theta)$$

Bernoulli distribution:

$$Y \sim \text{Bern}(\phi) \quad p(y | \phi) = \begin{cases} \phi, & y = 1 \\ 1 - \phi, & y = 0 \end{cases}$$

What is the likelihood for three i.i.d. samples, given parameter ϕ :

$$\mathcal{D} = \{y^{(1)} = 1, y^{(2)} = 1, y^{(3)} = 0\}$$

$$\begin{aligned} p(\mathcal{D} | \phi) &= \prod_{i=1}^N P(Y = y^{(i)} | \phi) \\ &= \phi \cdot \phi \cdot (1 - \phi) \end{aligned}$$

MLE

Estimating Parameters with Likelihood

We model the outcome of a single mysterious weighted-coin flip as a Bernoulli random variable:

$$Y \sim \text{Bern}(\phi)$$
$$p(y \mid \phi) = \begin{cases} \phi, & y = 1 \text{ (Heads)} \\ 1 - \phi, & y = 0 \text{ (Tails)} \end{cases}$$

Given the ordered sequence of coin flip outcomes:

$$[1, 0, 1, 1]$$

What is the estimate of parameter $\hat{\phi}$?

Estimating Parameters with Likelihood

We model the outcome of a single mysterious weighted-coin flip as a Bernoulli random variable:

$$Y \sim \text{Bern}(\phi)$$
$$p(y | \phi) = \begin{cases} \phi, & y = 1 \text{ (Heads)} \\ 1 - \phi, & y = 0 \text{ (Tails)} \end{cases}$$

Given the ordered sequence of coin flip outcomes:

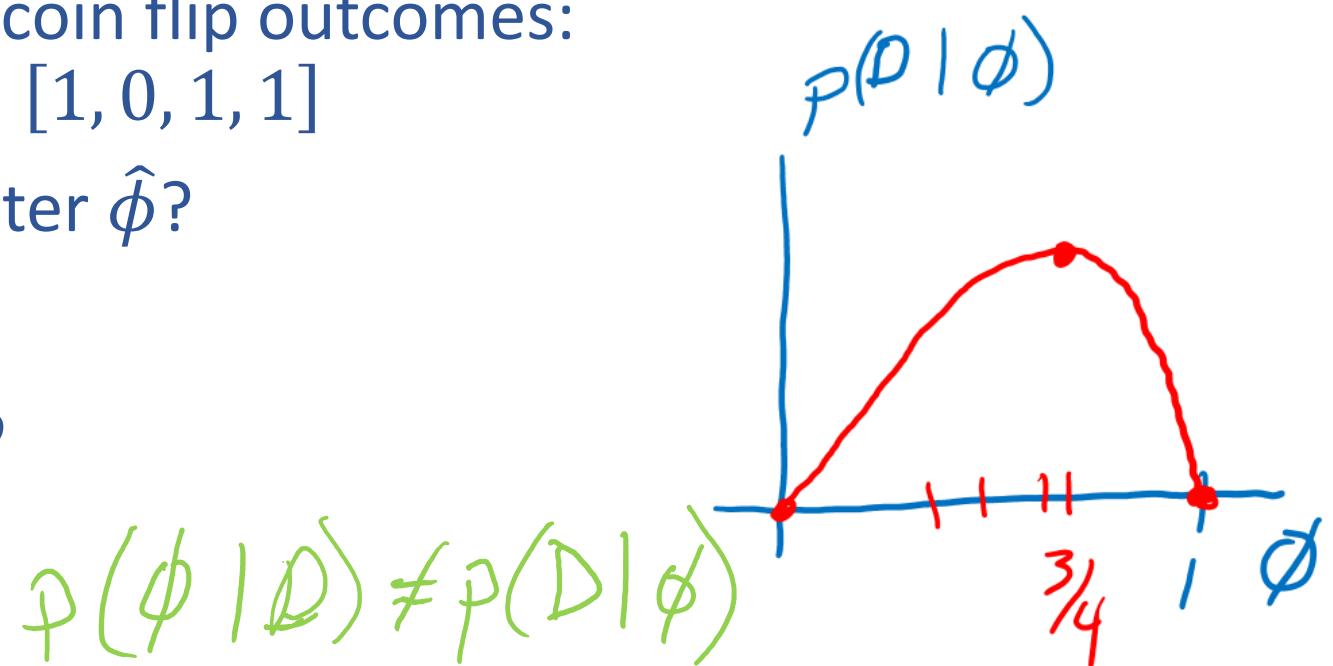
$$[1, 0, 1, 1]$$

What is the estimate of parameter $\hat{\phi}$?

$$\mathcal{L}(D | \phi) =$$

$$p(D | \phi) = \phi \cdot \phi \cdot (1 - \phi) \cdot \phi$$
$$= \phi^3 (1 - \phi)^1$$

$$p(\phi | D) \neq p(D | \phi)$$



Likelihood and Maximum Likelihood Estimation

Likelihood: The probability (or density) of random variable Y taking on value y given the distribution parameters, θ .

$$p(D|\theta)$$

Likelihood function: The value of likelihood as we change theta (same as likelihood, but conceptually we are considering many different values of the parameters)

$$L(\theta)$$

Maximum Likelihood Estimation (MLE): Find the parameter value that maximizes the likelihood.

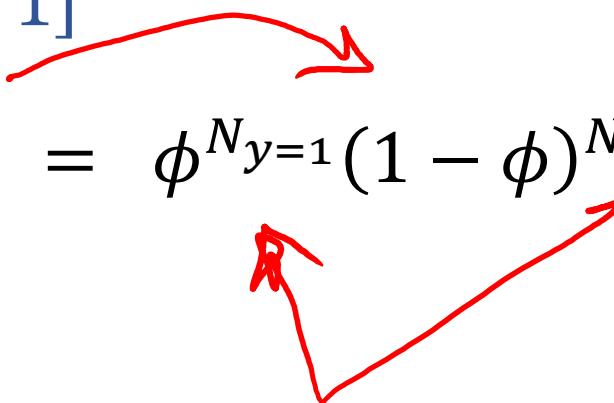
$$\hat{\theta}_{MLE} = \arg \max_{\theta} p(D|\theta) = \operatorname{argmax}_{\theta} L(\theta)$$

MLE as Data Increases

Given the ordered sequence of coin flip outcomes:

$$p(\mathcal{D} \mid \phi) = \prod_i^N p(y^{(i)} \mid \phi) = \phi^{N_{y=1}} (1 - \phi)^{N_{y=0}}$$

$[1, 0, 1, 1]$



What happens as we flip more coins?

MLE for Gaussian

Gaussian distribution:

$$Y \sim \mathcal{N}(\mu, \sigma^2)$$

$$\underline{p(y \mid \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y-\mu)^2}{2\sigma^2}}}$$

What is the log likelihood for three i.i.d. samples, given parameters μ, σ^2 ?

$$\mathcal{D} = \{y^{(1)} = 65, y^{(2)} = 95, y^{(3)} = 85\}$$

$$L(\mu, \sigma^2) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y^{(i)}-\mu)^2}{2\sigma^2}}$$

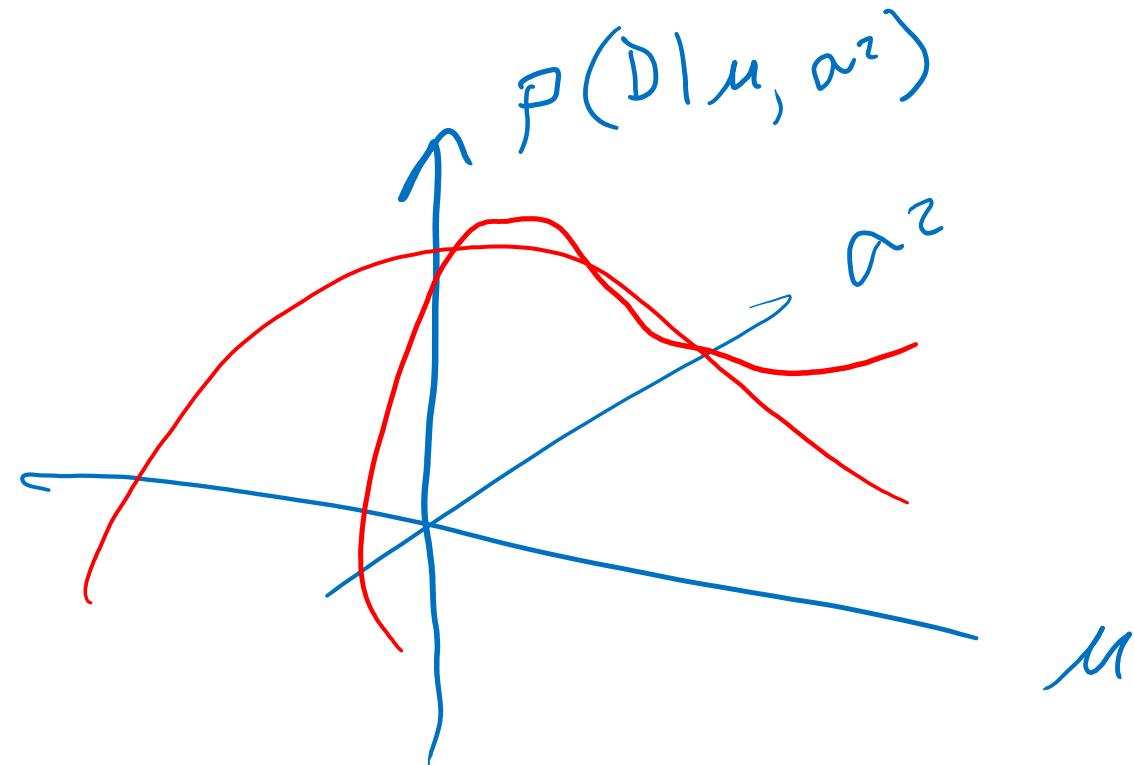
$$\hat{\theta}_{MLE} = \operatorname{argmax}_{\theta} \prod_i^N p(y^{(i)} \mid \theta)$$

MLE for Gaussian

Assume that exam scores are drawn independently from the same Gaussian (Normal) distribution.

Given three exam scores 75, 80, 90, which pair of parameters is the best fit (the highest likelihood)?

$$p(\mathcal{D}|\mu, \sigma^2) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y^{(i)}-\mu)^2}{2\sigma^2}}$$



MLE

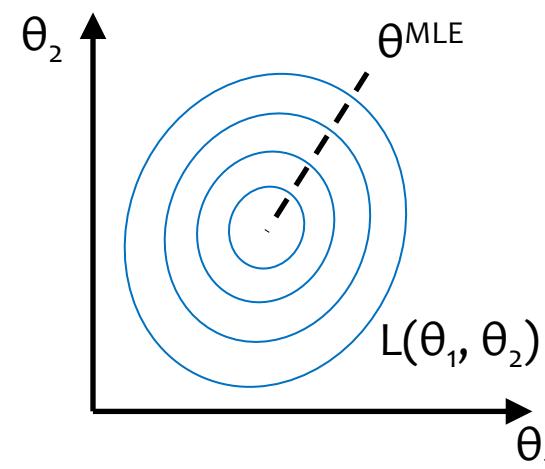
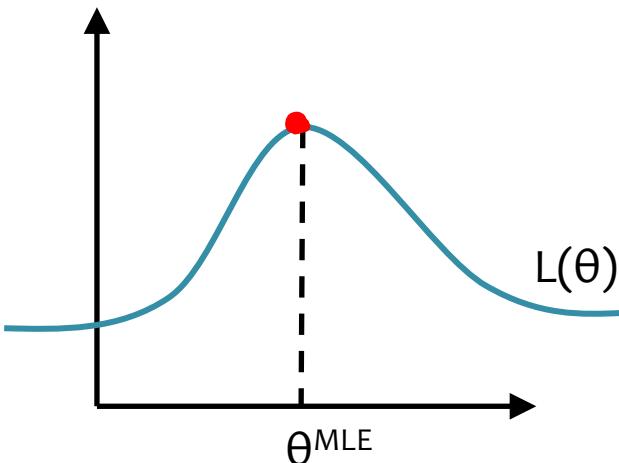
Suppose we have data $\mathcal{D} = \{x^{(i)}\}_{i=1}^N$

Principle of Maximum Likelihood Estimation:

Choose the parameters that maximize the likelihood of the data.

$$\theta^{\text{MLE}} = \underset{\theta}{\operatorname{argmax}} \prod_{i=1}^N p(\mathbf{x}^{(i)} | \theta)$$

Maximum Likelihood Estimate (MLE)



MLE Recipe

Likelihood and Log Likelihood

Likelihood: The probability (or density) of random variable Y taking on value y given the distribution parameters, θ .

$$P(D|\theta) = \prod p(y^{(i)}|\theta)$$

Likelihood function: The value of likelihood as we change theta (same as likelihood, but conceptually we are considering many different values of the parameters)

$$\text{likelihood } L(\theta; D) = P(D|\theta) = \prod p(y^{(i)}|\theta)$$

$$\text{log likelihood } \ell(\theta; D) = \log P(D|\theta) = \sum \log p(y^{(i)}|\theta)$$

$$\log xz = \log x + \log z$$

Maximum Likelihood Estimation

MLE of parameter θ for i.i.d. dataset $\mathcal{D} = \{y^{(i)}\}_{i=1}^N$

$$\hat{\theta}_{MLE} = \operatorname{argmax}_{\theta} p(\mathcal{D} | \theta)$$

$$= \operatorname{argmax}_{\theta} \prod_{i=1}^N p(y^{(i)} | \theta)$$

$$= \operatorname{argmax}_{\theta} \log \prod_{i=1}^N p(y^{(i)} | \theta)$$

$$= \operatorname{argmax}_{\theta} \sum_{i=1}^N \log p(y^{(i)} | \theta)$$

Recipe for Estimation

MLE

1. Formulate the likelihood, $p(\mathcal{D} \mid \theta)$
2. Set objective $J(\theta)$ equal to negative log likelihood

$$J(\theta) = -\log p(\mathcal{D} \mid \theta)$$

3. Compute derivative of objective, $\partial J / \partial \theta$
4. Find $\hat{\theta}$, either
 - a. Set derivate equal to zero and solve for θ
 - b. Use (stochastic) gradient descent to step towards better θ

MLE for Gaussian

Gaussian distribution:

$$Y \sim \mathcal{N}(\mu, \sigma^2)$$

$$p(y \mid \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y-\mu)^2}{2\sigma^2}}$$

What is the log likelihood for three i.i.d. samples, given parameters μ, σ^2 ?

$$\mathcal{D} = \{y^{(1)} = 65, y^{(2)} = 95, y^{(3)} = 85\}$$

$$L(\mu, \sigma^2) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y^{(i)}-\mu)^2}{2\sigma^2}}$$

$$\ell(\mu, \sigma^2) = \sum_{i=1}^N -\log \sqrt{2\pi\sigma^2} - \frac{(y^{(i)} - \mu)^2}{2\sigma^2}$$

$$\hat{\theta}_{MLE} = \operatorname{argmax}_{\theta} \prod_{i=1}^N p(y^{(i)} \mid \theta)$$

$$\hat{\theta}_{MLE} = \operatorname{argmax}_{\theta} \sum_{i=1}^N \log p(y^{(i)} \mid \theta)$$

Exercise: MLE for Exponential

Exponential distribution pdf:

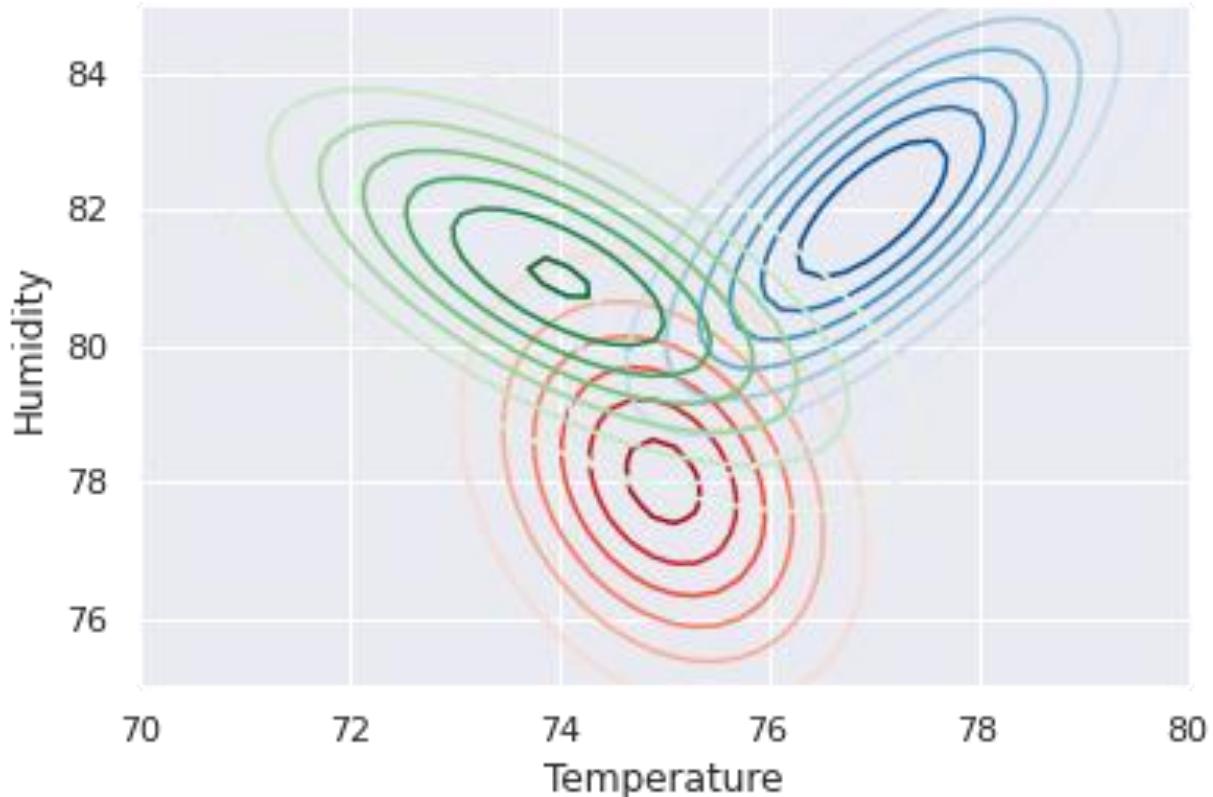
$$f(x) = \lambda e^{-\lambda x}$$

Two Applications of Bayes Rule

Prev Recitation

City classification

There are three nearby cities: **City A**, **City B**, and **City C**. Let $Y \sim \text{Categorical}(a, b, c)$ be a categorical distribution where $Y = 1$ means a randomly sampled sensor is in **City A**, $Y = 2$ means it is in **City B**, and $Y = 3$ means it is in **City C**.



Poll 2

Which of these terms is the likelihood?

$$p(y | x) = \frac{p(x | y) p(y)}{p(x)}$$

55 A →

30 B →

C →

D ↑

✓ 15% None of above

Bayes Rule

Terminology

Conditional
likelihood

Posterior

class
conditional class prior

$$p(y|x) = \frac{p(x|y)p(y)}{p(x)} = \frac{p(x,y)}{p(x)} = \frac{p(D)}{p(x)}$$

$$p(y|x, \theta) = \frac{p(x|y, \mu_y, \Sigma_y) p(y|q_1, q_2, q_3)}{p(x)}$$

Bayes Rule

Inserting parameters

$$p(y | x) = \frac{p(x | y) p(y)}{p(x)}$$

Bayes Rule

Another way to use Bayes rule

Posterior

Likelihood prior

$$p(\theta | x) = \frac{p(x | \theta) p(\theta)}{p(x)}$$

Poll 3

Where do we plug in the pdf, $f(x) = \lambda e^{-\lambda x}$

$$p(\theta | x) = \frac{p(x | \theta) p(\theta)}{p(x)}$$

A
B
C
D

MLE vs. MAP

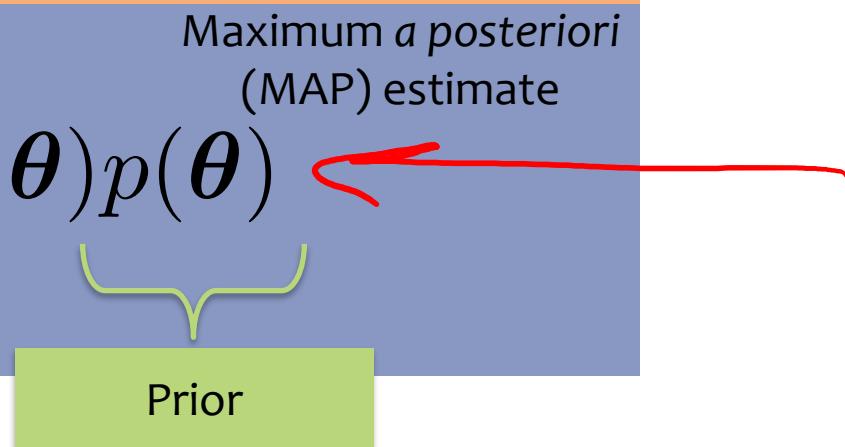
Suppose we have data $\mathcal{D} = \{x^{(i)}\}_{i=1}^N$

$$\theta^{\text{MLE}} = \operatorname{argmax}_{\theta} \prod_{i=1}^N p(\mathbf{x}^{(i)} | \theta)$$

Maximum Likelihood Estimate (MLE)

$$\theta^{\text{MAP}} = \operatorname{argmax}_{\theta} \prod_{i=1}^N p(\mathbf{x}^{(i)} | \theta) p(\theta)$$

Maximum *a posteriori* (MAP) estimate



Prior

Expectation and Variance

Expectation and Variance

The **expected value** of X is $E[X]$. Also called the mean.

- Discrete random variables:

Suppose X can take any value in the set \mathcal{X} .

$$E[X] = \sum_{x \in \mathcal{X}} xp(x)$$

- Continuous random variables:

$$E[X] = \int_{-\infty}^{+\infty} xf(x)dx$$

Expectation and Variance

The **variance** of X is $Var(X)$.

$$Var(X) = E[(X - E[X])^2]$$

- Discrete random variables:

$$Var(X) = \sum_{x \in \mathcal{X}} (x - \mu)^2 p(x)$$

$$\mu = E[X]$$

- Continuous random variables:

$$Var(X) = \int_{-\infty}^{+\infty} (x - \mu)^2 f(x) dx$$